

SAM (REX/Ops) Operations Meeting 2008.09.09

Adam (typist), Robert, Art, SteveS, SteveW, EricW

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Where did the summer go?

Releases:

Working on a new DB server. Recent problems with CDF DB's have exposed a potential vulnerability if commit and rollback fails [e.g. Oracle did not signal a rollback]. (this should not happen, but has happened in the CDF DB). When this happens dump the connection and make a new one. We want to prevent a subsequent commit from committing something unintentional (at this point Oracle must be rather sick).

The other DB server item - request from Herb to turn on single character wildcards. (e.g. use "_" - which we should not use). Robert wants to use "_". What about "?" - complaint against using unix and SQL wildcards together. Also, some of the SRM locations have literal ? in their names. Since \ sometimes does signal special meaning (like \n) then I guess _ is ok. Implementing this should be trivial. Do we have any files with _ in them? Probably not.

Eric: What about having the station ride out a long DB outage? Robert: has an implementation in the new Omniorb Station -- needs testing.

DØ:

Shifters' notes:

-- No SAM shifter Sept. 5-7 and most of Sept. 8 .

-- Sept. 3: At 17:36, there was a warning in the Raw Data Log (at

<http://d0om.fnal.gov/sam/DHAtAGlance/>):

"Wed Sep 3 17:41:09 2008 1220481669 Warning: d0ol-daq-4b -- There may be problems with dlsam or ENSTORE

Wed Sep 3 17:36:07 2008 1220481367 Action: Experts have been notified, please wait"

=====> It was OK later : "Wed Sep 3 17:46:05 2008 1220481965 OK"

Hardware problem with an online machine?

-- Sept. 4: There was some instability in the Mass storage status. At 2 pm the

D41FLTO.mover was in red status and the enstore overall status was yellow. Sent an e-mail to enstore-admin at 14:50. At 3 pm the mass

storage was in green status again. Stanley from the enstore-admin took care of the problem.

Also, some enstore problem on Friday that prevented production from getting files. Not very well known right now.

-- Sept. 7: "Station not responding" ---a major problem like last Monday. srm stager and samgfw02 had problem (reboot) about the same time twice.

Have a helpdesk ticket open to look at samgfw02's logs -- Steve Timm looked and didn't see anything.


Lots of CORBA errors last night - station crashed and restarted at 5am. Still looking at it. Purdue SRM stager restarted just before midnight.

Have a question into Andrew asking why the Purdue stager is causing problems. Mike asks to disable Purdue.

Adam's new SAMGrid tracking.

See <http://d0db-prd.fnal.gov/rexipedia/common/SAMGridD0/prodEff>

Monday 2008 September 1 - Monday 2008 September 8



1. **Monday (labor day):** See last week's report below.
2. **Tuesday:** See last week's report below.
3. **Tuesday - Saturday:** Note period of very good running, especially starting on Friday. Note the large number of queued jobs. That was when Heidi Schellman started submitting jobs, so we had both Mike and Heidi feeding the queue. Also, we think that Heidi gets her own 750 job/site limit, which keeps more jobs in the system.
4. **Sunday:** See [DZSAM-244](#). This problem was a reprise of the SRM blocking the SAM station that we saw last week (see below). Curiously, samgfw02 (forwarding node #2) spontaneously rebooted during this time, just like last week.

Recent unresolved issues:

DZSAM-235: "Lancaster-LCG not visible on SAMGRID monitoring page" : Willem sent his newest problem

Had been working on last of the issues - authentication issue at the SRM at Imperial (we now think this is a firewall between catapiller and lancaster - tests on catapiller fail, but on d0mino it works - it's a firewall. Sabah). Jobs were getting there and output was being returned. Last week power outage at Lancaster. Now jobs that get submitted fail because of RSL retries (RSL = global resource manager). Willem's jobs are being aborted. Phone meeting today. Edg job submit command fails to load one of the libraries? Was edg updated?

DZSAM-236: User "Herb" asks about single character wildcard. Have we decided to do anything about this ?

DZSAM-237: "Low space on samgfw01 sandbox" : so what do we do ?

DZSAM-238 : "Forwarding nodes submit jobs to CAB very slowly" : Adam created this one himself.

What about problems at Fermigrid where it could not accept jobs quickly? There were bluearc problems last week.

DZSAM-241: "slow/no file delivery to luhep" : It seems that Robert had solved the problem. Should we close this ?

Close it - Joel had two different stagers battling over the disk

DZSAM-244: "Station not responding" (as mentioned above in the Shifters' Notes)

SRM problem

DZSAM-246: "Does dbserver_/d0ora2/_dbs_grid1_prd correspond to jobs running from

samgfw01 and same for 2?" : Can someone answer ?

Heidi's question - see ticket.

CDF:

CDF SAM/DH Shift Report 09/02/08 - 09/08/08

very quiet shift, most noticeable problems:

Sep. 2nd: fcdldata320 offline (from previous shift): back online: fixed.

Just had to reinstall the dCache software.

Sep. 2nd: b0csl21 does not allow log-in (machine pingable), request sent to helpdesk (by Angela): fixed. *By rebooting*

Sep. 3rd: CDFDH-244: slow IO rate noticed from users on some dcache pool, the restart of the dCache services on these

pools seems to have fixed the "hot" pool problem (Angela): waiting from user feedback to close the ticket. *Angela drained the pools, but still hot.*

Sep. 3rd: CDFDH-245: errors from upload server 322: fixed.

Sep. 4th: CDFDH-250: fcdldata109 OFFLINE, rebooted (Etta): fixed.

Sep. 4th: Tape IAB920 on CDF-LTO3 marked NOACCESS on DH at a glance and ENCP history page: fixed.

DB connection problems caused file deliveries to stop (max session error) due to commits and rollbacks were failing. Were getting partial commits - this should NEVER happen. Oracle admits this is an "unpublished bug" (a bug they don't publicize). The workaround is to increase the number of sessions and that will happen for the next downtime. Second time we had the partial commit it corrupted an index (!!!). That prevented us from fixing the inconsistencies. DBAs had to drop the index and recreate to fix.

MINOS:

DB:

See CDF above under Ray.

Got the sam_read # session exceeded error (BUT THIS MAY BE MISLEADING - SEE BELOW). But also got trigger errors.

Program noted by Oracle is /fbsng/condor_local/execute/dir_24518/work/bin (just the local directory for the executable). /local/stage1/condor/dir_nnnnn/work/bin .

Can see the queries. Nelly will look. Performance on cdfopr2 is slow because of this.

An example query is...

```
SELECT /*+ RULE; */ SUM(LUM_INTEGRAL_ONLINE) ,  
SUM(LUM_INTEGRAL_OFFLINE) FROM DATA_FILES JOIN  
DATA_FILES_LUMBLOCKS USING(FILE_ID) JOIN  
FILECATALOG.CDF2_RUNSECTIONS ON( ID BETWEEN LUM_MIN  
AND LUM_MAX) WHERE FILE_NAME = 'er040143.02eaphys'
```

Looks like a production query.

Also a trigger error for sam_prod_ap_update (??). - Maybe this was a connection problem.

User I/O on the machine is high, generating some waiting.

Right now waitstatus of 6 and idle 25%; normal is waitstatus of 2, idle of 40%

No communication with Ray yet.

***** BUT - Max sessions exceeded error may be misleading because it happens on commit when you already have a session. This may be a symptom of some other Oracle problem and Oracle just spits out this error. But there are still a lot of samread sessions.*

Enstore/dCache:

No report

Special: Power outage on 9/18 at GCC

- Thursday, September 18th -- GCC Power Outage
- 480V repair, testing, maintenance - **no generators**
- CRA 1-3 hours, CRB 2-8 hours (CRA = computer room A)
- will bring cooling switch boards back first
- expect GCC networking & tape libraries to remain up on UPS
- GCC/CRC and LCC 480 maintenance tbd (Sept. 17, 18, 19)
- try to schedule LCC 13kV maintenance with LCC 480V maintenance
- date still tentative based on availability of FESS resources